Optimally Influencing Complex Ising Systems

Christopher Lynn[1] and Daniel D. Lee[2]

[1]*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*
[2]*Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

arXiv:1609.09180v1 [cond-mat.dis-nn] 29 Sep 2016

In the study of social networks, a fundamental problem is that of influence maximization (IM): How can we maximize the collective opinion of individuals in a network given constrained marketing resources? Traditionally, the IM problem has been studied in the context of contagion models, which treat opinions as irreversible viruses that propagate through the network. To study reverberant opinion dynamics, which yield complex macroscopic behavior, the IM problem has recently been proposed in the context of the Ising model of opinion dynamics, in which individual opinions are treated as spins in an Ising system. In this paper, we are among the first to explore the *Ising influence maximization (IIM)* problem, which has a natural physical interpretation as the maximization of the magnetization given a budget of external magnetic field, and we are the first to consider the IIM problem in general Ising systems with negative couplings and negative external fields. For a general Ising system, we show analytically that the optimal external field (i.e., that which maximizes the magnetization) exhibits a phase shift from intuitively focusing on high-degree nodes at high temperatures to counterintuitively focusing on "loosely-connected" nodes, which are weakly energetically bound to the ground state, at low temperatures. We also present a novel and efficient algorithm for solving IIM with provable performance guarantees for ferromagnetic systems in nonnegative external fields. We apply our algorithm on large random and real-world networks, verifying the existence of phase shifts in the optimal external fields and comparing the performance of our algorithm with the state-of-the-art mean-field-based algorithm.

I. INTRODUCTION

With the proliferation of online social networks over the last two decades, the problem of optimally influencing the opinions of individuals in social networks has garnered much attention [8, 19, 36]. The prevailing paradigm treats marketing as a viral process and is studied in the context of contagion models and irreversible processes [7, 15, 19, 28]. Under the viral paradigm, the influence maximization (IM) problem is, given a budget of seed infections, to choose the subset of individuals to infect such that the ensuing cascade is maximized at the completion of the viral process. While the viral paradigm accurately describes out-of-equilibrium phenomena, such as the introduction of new ideas or products to a system, this paradigm fails to capture reverberant opinion dynamics wherein repeated interactions between individuals in the network give rise to complex macroscopic opinion patterns, as, for example, is the case in the formation of political opinions [12, 17, 25, 29, 38]. In such cases, rather than maximizing the spread of a viral advertisement, the marketer is interested in optimally shifting the equilibrium opinions of individuals in the network.

To study complex opinion formation, the influence maximization problem has recently been proposed in the context of the Ising model of opinion dynamics, in which individual opinions are treated as spins in an Ising system at dynamic equilibrium and marketing to the system is modeled as the addition of an external magnetic field [23, 24]. The resulting problem, known as the *Ising influence maximization (IIM)* problem, is equivalent to the maximization of the magnetization of an Ising system given a budget of external field. Figure 1 gives an intuition for the IIM problem in a small preferential attachment network. In Figure 1(a) we find that, compared with the trivial strategy of spreading the external field uniformly among the nodes, focusing the external field on influential nodes can increase the magnetization of the system by up to 60%. Remarkably, we find in Figure 1(b) that the structure of the optimal external field (i.e. that which maximizes the magnetization) undergoes a phase shift: at high temperatures, the magnetization is maximized by intuitively focusing on the single node of highest degree, while at low temperatures, the magnetization is maximized by counterintuitively spreading the external field among nodes of low degree.

In this paper, we explore the structure of optimal external fields in complex Ising systems, and we provide novel analytic and algorithmic tools for identifying influential nodes in large complex networks. One might suspect that the dramatic phase shift seen in Figure 1(b) is limited to a small class of Ising systems, and while previous work on the IIM problem has been limited to ferromagnetic systems in nonnegative external fields, we show analytically that such phase shifts are a persistent phenomenon found in all Ising systems, including those with negative couplings and negative external fields. In particular, for general Ising systems, we provide analytic solutions to IIM in the high-T and low-T limits, showing that the optimal external fields undergo a phase shift from focusing on high-degree nodes at high temperatures to focusing on "loosely-connected" nodes, which are weakly energetically bound to the ground state, at low temperatures.

We are also the first to study the Ising influence maximization problem directly (without mean-field approximations), and we provide a novel and efficient algorithm for solving IIM with provable performance guarantees for ferromagnetic systems in nonnegative external fields. We

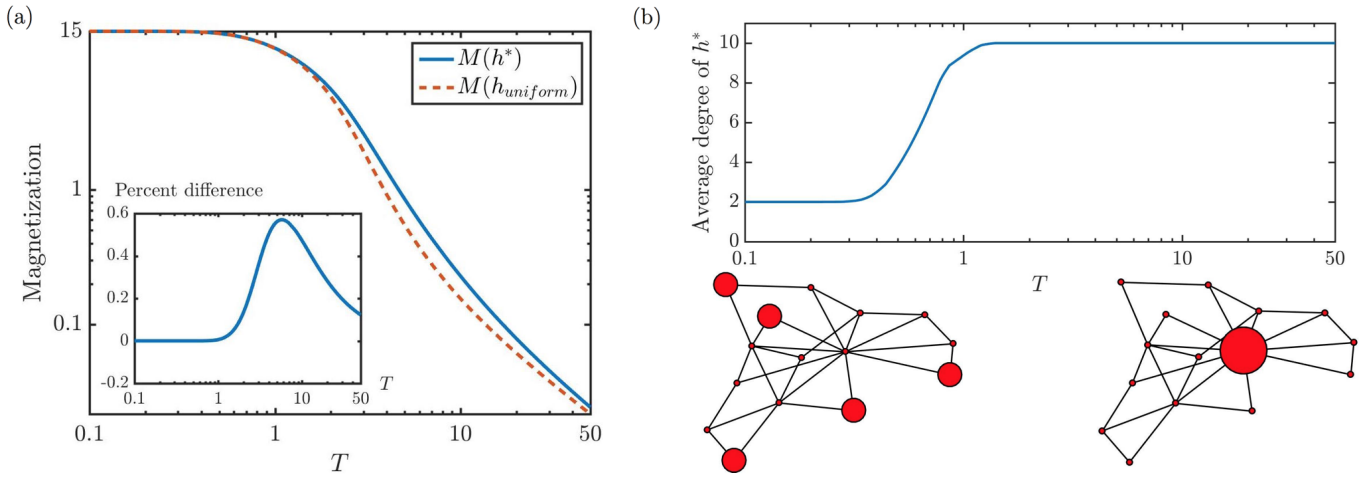

FIG. 1. We consider a preferential attachment network consisting of $n = 15$ nodes with ferromagnetic bonds of weight 1 and no initial external field. For a budget $H = 1$, we can apply any additional external field $\boldsymbol{h}$ subject to the constraint $|\boldsymbol{h}|_1 \leq H$. (a) The magnetization of the system under the optimal external field $\boldsymbol{h}^*$ is compared with the magnetization under $\boldsymbol{h}_{uniform} = \frac{H}{n}(1, \ldots, 1)^T$, revealing that the optimal external field yields up to 60% larger magnetization. (b) The optimal external field undergoes a phase shift from spreading H among the nodes of lowest degree at low temperatures to focusing on the node of highest degree at high temperatures.

apply our algorithm on large random and real-world networks, verifying the existence of a phase shift in the optimal external fields and comparing the performance of our algorithm with the state-of-the-art mean-field-based algorithm in [24].

The paper is organized as follows. In Section II we introduce the Ising model of opinion dynamics and formally state the IIM problem. In Sections III and IV we derive analytic solutions to IIM in the the high-T and low-T limits, respectively. In Section V we present an efficient algorithm for solving IIM, and we prove performance guarantees for the case of ferromagnetic systems in nonnegative external fields. In Section VI we provide numerical experiments on large random and real-world networks, verifying the existence of phase shifts in the optimal external fields and comparing our algorithm with the state-of-the-art. We conclude with summarizing remarks in Section VII.

I.1. Related work

The recent development of the influence maximization problem can be traced to Domingos and Richardson [8, 36] and Kempe et al. [19], the latter of which presented the problem in the context of the independent cascade model, marking the introduction of the viral paradigm that is still widely used today. Subsequent work has expanded the IM problem to include competing viruses [15], fractional initial infections [7], and continuous-time dynamics [9, 14], among many other interesting generalizations. However, all of this work is fundamentally based on contagion models, which do not allow for reverberant opinion dynamics and complex steady-states.

Our goal is to study the IM problem in contexts where individuals flip opinions as they interact, yielding complex collective behavior, and the Ising model is a natural starting point given its widespread success in describing complex systems. The Ising model has been applied to opinion formation dating to the work of Galam [11, 13], and continues to be an active area of research [6, 27]. Ising-like models are also prevalent under names such as logistic response in game theory [4, 26] and the Sznajd model in sociophysics [39]. Furthermore, complex Ising models have found widespread use in machine learning, and so the model we consider is formally equivalent to a pair-wise Markov random field or a Boltzmann machine [33, 40, 42].

The Ising influence maximization problem represents a paradigm shift in the study of influence maximization and bridges a gap between the IM and Ising literatures. The IIM problem was formalized in [24], where the authors limited their study to ferromagnetic systems in nonnegative external fields and only considered the IIM problem under the mean-field approximation. Our work represents two large steps forward: (i) we consider the IIM problem in general networks under general external fields, and (ii) we study the IIM problem directly via exact Ising calculations for small networks and Monte Carlo simulations for large networks. In order connect our results with previous work, we also develop a number of analytic results under the variational mean-field approximation, which has its roots in statistical mechanics and has long been used for inference in machine learning [18, 35, 37]. We note that Liu et al. also studied a similar problem in [23], but they neglected the effects of thermal noise, which we find to play a vital role.

II. ISING INFLUENCE MAXIMIZATION

In this section, we present the Ising model of opinion dynamics and describe how negative couplings and negative external fields account for two important features found in real-world social networks: (i) negatively correlated opinions and (ii) competitive marketing. We conclude by formally stating the Ising influence maximization problem in this generalized context.

II.1. Generalized Ising model of opinion dynamics

We consider an Ising system consisting of a set of individuals $N = \{1, \dots, n\}$, each of which is assigned an opinion $\sigma_i \in \{\pm 1\}$ that captures its current state. By analogy with the Ising model in statistical physics, we often refer to individuals as nodes and opinions as spins, and we refer to $\boldsymbol{\sigma} = (\sigma_i)$ as a spin configuration of the system. Individuals in the network interact via a symmetric coupling matrix J, where $J_{ij} = J_{ji} \in \mathbb{R}$, and we assume the diagonal entries in J are zero such that self-interactions are forbidden. Each node also interacts with forces external to the network via an external field $\boldsymbol{h} = (h_i)$, where $h_i \in \mathbb{R}$ is the external field on node i. For example, if the spins represent the political opinions of individuals in a social network, then J_{ij} represents the the amount of influence that individuals i and j hold over each other's opinions and h_i represents the political bias of node i due to external forces such as campaign advertisements and news articles.

We note that $\boldsymbol{h}$ and J are allowed both positive and negative entries of arbitrary weight, whereas previous models have only considered $\boldsymbol{h} \geq 0$ and $J \geq 0$. A negative external field $h_i < 0$ represents a competitive external influence on node i. Competing sources of external influence are natural in most marketing contexts and have been studied extensively in the influence maximization literature [3, 15]. A negative coupling $J_{ij} < 0$ signifies that the spins i and j are anticorrelated, such that if individual i has a positive opinion, then individual j will be influenced to have a negative opinion, and vice versa. Anticorrelated opinions have been studied in the psychology literature [34], but, to the authors' knowledge, have not yet been considered in models of influence maximization, even in the viral paradigm. Finally, we note that the introduction of negative couplings makes our model similar to a spin-glass, which is known to show much more complex behavior than a standard Ising ferromagnet [32].

The steady-state of an Ising system is described by the Boltzmann distribution

$$P(\boldsymbol{\sigma}) = \frac{1}{Z} \exp\left(-\frac{1}{T} E(\boldsymbol{\sigma})\right), \tag{1}$$

where T is the temperature of the system, $E(\boldsymbol{\sigma}) = -\frac{1}{2}\sum_{ij} J_{ij}\sigma_i\sigma_j - \sum_i h_i\sigma_i$ is the energy of configuration $\boldsymbol{\sigma}$, and $Z = \sum_{\{\boldsymbol{\sigma}\}} \exp\left(-\frac{1}{T}E(\boldsymbol{\sigma})\right)$ is the partition function. Unless otherwise specified, sums over i and j run over the set of nodes N, and sums over $\{\boldsymbol{\sigma}\}$ are assumed over the set of spin configurations $\Omega = \{\pm 1\}^n$. The expectation of any function $f(\boldsymbol{\sigma})$ over (1) is denoted by $\langle f(\boldsymbol{\sigma})\rangle = \sum_{\{\boldsymbol{\sigma}\}} f(\boldsymbol{\sigma})P(\boldsymbol{\sigma})$. In particular, the total expected spin $M = \sum_i \langle\sigma_i\rangle$ is referred to as the *magnetization* of the system, and we will often consider the magnetization as a function of the external field, denoted by $M(\boldsymbol{h})$.

Another important concept is the *susceptibility*, represented by the matrix

$$\chi_{ij} = \frac{\partial \langle\sigma_i\rangle}{\partial h_j} = \frac{1}{T}\left(\langle\sigma_i\sigma_j\rangle - \langle\sigma_i\rangle\langle\sigma_j\rangle\right), \tag{2}$$

where χ_{ij} quantifies the response of node i to a change in the external field on node j. Identifying the susceptibility matrix with the spin covariance matrix, $\chi_{ij} = \frac{1}{T}\mathrm{Covar}(\sigma_i, \sigma_j)$, it follows that χ is symmetric. Furthermore, to illustrate the importance of χ, we note that, given vanishingly small external field resources δH, the magnetization is maximized by focusing δH on the nodes corresponding to the largest elements of the gradient $\frac{\partial M}{\partial h_i} = \sum_j \chi_{ji}$. This intuition will prove useful in Section III.

Notation. Unless otherwise specified, capital letters represent matrices and bold symbols represent column vectors with the appropriate number of components, while non-bold symbols with subscripts represent individual components. We often abuse notation and write relations such as $\boldsymbol{m} \geq 0$ to mean $m_i \geq 0$ for all components i. Along the same lines, any descriptor such as "nonnegative," when applied to a vector or matrix, is assumed to be applied component-wise.

II.2. Problem formulation

We study the problem of maximizing the total magnetization of an Ising system with respect to the external field. For concreteness, we assume that an external agent can impose an external field $\boldsymbol{h}$ on the network subject to the constraint $|\boldsymbol{h}|_1 \leq H$, where $|\boldsymbol{h}|_1 = \sum_i |h_i|$ is the ℓ_1 norm and $H > 0$ is the *external field budget*, and we denote the set of feasible external fields by $\mathcal{F}_H = \{\boldsymbol{h} \in \mathbb{R}^n : |\boldsymbol{h}|_1 \leq H\}$. In general, we also assume that the system experiences an *initial external field* $\boldsymbol{h}^{(0)}$, which cannot be controlled by the external agent. We are now prepared to state the Ising influence maximization problem.

Problem 1 (*Ising influence maximization*)**.** Given a system described by J, $\boldsymbol{h}^0$, and T, and an external field budget H, find an external field

$$\boldsymbol{h}^* = \operatorname*{arg\,max}_{\boldsymbol{h}\in\mathcal{F}_H} M(\boldsymbol{h}^0 + \boldsymbol{h}). \tag{3}$$

We refer to any such external field as optimal. We also often refer to any nodes included in an optimal external

field as optimal.

Remark. Before proceeding, we point out that for finite systems, the objective $M(\boldsymbol{h}^0+\boldsymbol{h})$ is uniquely-defined for all $\boldsymbol{h}$ and hence the IIM problem is well-posed.

III. HIGH-T: HIGH-DEGREE NODES ARE OPTIMAL

In this section, we develop an analytic solution to IIM in the high-temperature limit. For a general Ising system described by J and $\boldsymbol{h}^{(0)}$ and an external field budget H, we prove in the $T\to\infty$ limit that the optimal external fields focus on the nodes with the highest degree in the network. While focusing on the nodes of highest degree is intuitive in that we tend to think of high-degree nodes as influential in driving network dynamics, we note that the independence of this result with respect to $\boldsymbol{h}^{(0)}$ is quite remarkable. We conclude by drawing comparisons with the external fields that maximize the mean-field magnetization, which have been the focus of previous work [24].

III.1. High-T analytic solution to IIM

In this subsection, we show that in the high-temperature limit, for a general Ising system and external field budget H, the magnetization is maximized by focusing H on the nodes of highest degree. In what follows we consider a general Ising system defined by J, $\boldsymbol{h}$, and T. Instead of studying the magnetization directly, we find it most intuitive to consider the susceptibility χ. In the high-T limit, we can expand $\frac{1}{\beta}\chi_{ij}=\langle\sigma_i\sigma_j\rangle-\langle\sigma_i\rangle\langle\sigma_j\rangle$ in powers of $\beta\equiv\frac{1}{T}$,

$$\frac{1}{\beta}\chi_{ij}=\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0}+\beta\frac{\partial}{\partial\beta}\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0}+\dots. \tag{4}$$

Since the Boltzmann distribution is uniform at $\beta=0$, we have $\langle f(\boldsymbol{\sigma})\rangle_{\beta=0}=\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}f(\boldsymbol{\sigma})$ for any function of the spins $f(\boldsymbol{\sigma})$. Thus, when evaluated at $\beta=0$, any terms in $f(\boldsymbol{\sigma})$ that include an odd number of spins will vanish in expectation. In particular, $\langle\sigma_i\rangle_{\beta=0}=0$ for all $i\in N$, and hence the zeroth-order term in Eq. (4) takes the form

$$\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0}=\langle\sigma_i\sigma_j\rangle_{\beta=0}=\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\sigma_i\sigma_j=\delta_{ij}. \tag{5}$$

In Appendix A, we calculate the first- and second-order terms in Eq. (4), yielding

$$\begin{aligned}\chi_{ij}=\;&\beta\delta_{ij}\\&+\beta^2J_{ij}\\&+\beta^3\left(\sum_k J_{ik}J_{kj}-\delta_{ij}\sum_k J_{jk}^2-\delta_{ij}h_i^2\right)\\&+\dots.\end{aligned} \tag{6}$$

Using Eq. (6), we can relate changes in the external field $\boldsymbol{h}$ to changes in the total magnetization M:

$$\frac{\partial M}{\partial h_i}\equiv\sum_j\chi_{ij}=\beta+\beta^2d_i+\beta^3\left(\sum_{j\neq i}(J^2)_{ij}-h_i^2\right)+\dots, \tag{7}$$

where $d_i=\sum_j J_{ij}$ is the *degree* of node i and $\sum_{j\neq i}(J^2)_{ij}=\sum_{jk}J_{ik}J_{kj}-\sum_j J_{ij}^2$ represents the total weight of paths of length 2 originating from node i (not counting self-interactions).

Inspecting Eq. (7), we find to first-order in β that $\frac{\partial M}{\partial h_i}>0$, which tells us that in the high-T limit, the optimal external fields $\boldsymbol{h}^*$ have positive components summing to H, i.e., $\sum_i h_i^*=H$. This may seem obvious, but we note that at intermediate and low temperatures, applying negative external field to nodes that are anti-correlated with the total magnetization can be optimal. To understand the dependence of the optimal external fields on the topology of the network, we turn to the second-order term in Eq. (7), which is proportional to the degree of node i. Since $\boldsymbol{h}^*\geq 0$ for all high-T optimal external fields, we conclude that in the high-T limit, the magnetization is maximized by focusing H on the nodes of highest degree in the network, independent of the initial external field. The result is summarized in Theorem 2.

Theorem 2. Let J and $\boldsymbol{h}^{(0)}$ describe a general Ising system and consider a general external field budget H. For sufficiently high temperatures, all optimal external fields place positive weight on the nodes of highest degree summing to H; i.e. letting $N_{max}=\arg\max_i d_i$ denote the set of nodes of maximum degree, if $\boldsymbol{h}^*\in\mathcal{F}_H$ maximizes $M(\boldsymbol{h}^{(0)}+\boldsymbol{h})$ in the high-T limit, then $\sum_{i\in N_{max}}h_i^*=H$, and hence $h_i^*=0$ for all $i\notin N_{max}$.

Focusing one's budget on the nodes of highest degree is intuitive since we tend to think of high-degree nodes as influential in driving network dynamics. However, the independence of this result with respect to $\boldsymbol{h}^{(0)}$ can be counterintuitive, especially if one considers the extreme case of a network in which the node of highest degree experiences a large initial external field.

If there exist multiple nodes of highest degree, then we turn to the third-order term in Eq. (7) to break the degeneracy between possible optimal external fields. The third-order term in Eq. (7) has a quadratic dependence on $\boldsymbol{h}$, favoring nodes with initial external fields that are small in magnitude. However, if $\boldsymbol{h}^{(0)}$ is uniform, then the high-T optimal external fields focus H on the nodes of highest degree which reach the largest number of nodes via paths of length 2. In the next subsection, we show that the mean-field magnetization is also maximized by focusing H on the nodes of highest degree, but that the third-order term in Eq. (7) is different under the mean-field approximation, leading to the possibility of sub-optimal external fields at high temperatures.

III.2. High-T analytic solution under the mean-field approximation

We now compare the optimal external fields to the mean-field optimal external fields (those which maximize the mean-field magnetization) in the high-temperature limit. We prove for general Ising systems that the high-T mean-field magnetization is maximized by focusing on the nodes of highest degree, formally verifying a conjecture made in [24]. Furthermore, we go on to show that the third-order term in the high-T expansion of the mean-field susceptibility is different from that in Eq. (7), leading to the possibility of sub-optimal external fields under the mean-field approximation.

For an Ising system described by J, $\boldsymbol{h}$, and T, the mean-field magnetization is defined by the self-consistency equations [18, 35, 42]:

$$m_i = \tanh\left[\beta\left(\sum_j J_{ij}m_j + h_i\right)\right], \tag{8}$$

for all $i \in N$. To calculate the high-T magnetization, we expand Eq. (8) with respect to β:

$$m_i = \beta\left(\sum_j J_{ij}m_j + h_i\right) - \frac{\beta^3}{3}\left(\sum_j J_{ij}m_j + h_i\right)^3 + \dots. \tag{9}$$

Because the leading-order term in $\boldsymbol{m}$ is first-order in β, we can drop $\sum_j J_{ij}m_j$ in the cubic term in Eq. (9). Abusing notation to let $\boldsymbol{h}^3 \equiv (h_1^3, \dots, h_n^3)^T$, we have

$$\begin{aligned} &\boldsymbol{m} = \beta(J\boldsymbol{m} + \boldsymbol{h}) - \frac{\beta^3}{3}\boldsymbol{h}^3 + \dots \\ \Rightarrow \quad &\boldsymbol{m} = (I - \beta J)^{-1}\left(\beta\boldsymbol{h} - \frac{\beta^3}{3}\boldsymbol{h}^3\right) + \dots, \end{aligned} \tag{10}$$

where I is the $n \times n$ identity matrix. Assuming that our solution to Eq. (8) is stable, i.e., $\boldsymbol{m}$ describes a minimum of the mean-field Gibbs free energy, we have that $\rho(\beta J) < 1$, where $\rho(\cdot)$ denotes spectral radius (see [24]). This allows us to expand $(I - \beta J)^{-1}$ as a geometric series of matrices, yielding

$$\boldsymbol{m} = \beta\boldsymbol{h} + \beta^2 J\boldsymbol{h} + \beta^3\left(J^2\boldsymbol{h} - \frac{1}{3}\boldsymbol{h}^3\right) + O(\beta^4). \tag{11}$$

From Eq. (11), one can easily derive the mean-field susceptibility of the system, which takes the form

$$\frac{\partial M_{MF}}{\partial h_i} = \beta + \beta^2 d_i + \beta^3\left(\sum_j (J^2)_{ij} - h_i^2\right) + \dots, \tag{12}$$

Eq. (12) verifies, for a general Ising system and external field budget H, that under the mean-field approximation the high-T optimal external fields are positive and correctly focus H on the nodes of highest degree. However, the third-order term in Eq. (12) differs from the third-order term in Eq. (7) in that the sum over $(J^2)_{ij}$ runs over all j instead of running over $j \neq i$. Thus, the mean-field approximation mistakenly takes into account self-interactions via paths of length 2, which can lead to sub-optimal external fields in the high-T limit, as shown in Figure 2.

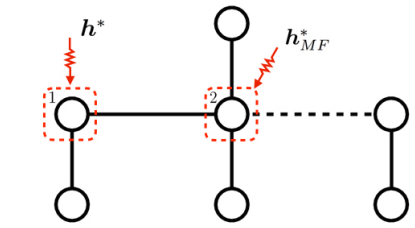

	d_i	$\sum_{j\neq i}(J^2)_{ij}$	$\sum_j (J^2)_{ij}$
node 1	2	1	3
node 2	2	0	4

ferromagnetic bond (weight 1) — highest degree
antiferromagnetic bond (weight -1) — positive external field

FIG. 2. A small network in the high-temperature limit with $J_{ij} \in \{\pm 1\}$ corresponding to ferromagnetic and antiferromagnetic bonds, respectively. To second-order in β, nodes 1 and 2 are optimal, independent of $\boldsymbol{h}^{(0)}$ and H, since they both have maximum degree. However, the third-order term breaks the degeneracy: for $\boldsymbol{h}^{(0)} = 0$ and $H < 1$, the exact magnetization is maximized by focusing H on node 1, while the mean-field magnetization is maximized by focusing H on node 2.

IV. LOW-T: "LOOSELY-CONNECTED" NODES ARE OPTIMAL

In this section, we study the structure of optimal external fields in the low-temperature limit. At low temperatures, the magnetization is governed by the spin configurations associated with low energies, i.e., the ground states, first-excited states, etc. However, for a large external field budget H, the feasible external fields $\boldsymbol{h} \in \mathcal{F}_H$ can alter the low-energy spectrum dramatically, making a general description of the low-temperature magnetization both cumbersome and unilluminating. To ease the discussion and foster intuition, we consider a natural special case: when the system admits a unique ground state and a unique first-excited state for all external fields $\boldsymbol{h} \in \mathcal{F}_H$. In this case, we show that the low-T magnetization is maximized by focusing H on the nodes with opposite parity between the ground state and first-excited state configurations, and we refer to such nodes as "loosely-connected" since they are weakly energetically bound to the ground state. For an analytic discussion of low-temperature solutions to IIM for general Ising systems and general external field budgets we point the reader to Appendix B. In order to connect our results to previous work that has focused on maximizing the mean-field magnetization, we conclude this section by drawing

comparisons with the low-T mean-field optimal external fields.

IV.1. A low-T analytic solution to IIM

In this subsection, we show that in the low-temperature limit, for an Ising system that admits a unique ground state and a unique first-excited state for all $\boldsymbol{h} \in \mathcal{F}_H$, the magnetization is maximized by focusing H on the nodes with opposite parity between the ground and first-excited states. We consider an Ising system described by J, $\boldsymbol{h}^{(0)}$, and T and an external field budget H such that the system admits unique ground and first-excited states $\boldsymbol{\sigma}^0$ and $\boldsymbol{\sigma}^1$, respectively, for all feasible external fields $\boldsymbol{h} \in \mathcal{F}_H$. In other words, we have that $\boldsymbol{\sigma}^0 = \arg\min_{\boldsymbol{\sigma}\in\{\pm 1\}^n} E_{\boldsymbol{h}}(\boldsymbol{\sigma})$ and $\boldsymbol{\sigma}^1 = \arg\min_{\boldsymbol{\sigma}\in\{\pm 1\}^n/\boldsymbol{\sigma}^0} E_{\boldsymbol{h}}(\boldsymbol{\sigma})$ for all $\boldsymbol{h} \in \mathcal{F}_H$, where $E_{\boldsymbol{h}}(\cdot)$ denotes the energy under the external field $\boldsymbol{h}^{(0)} + \boldsymbol{h}$. To give a sense of the systems for which this condition holds, we note that any system that initially (under $\boldsymbol{h}^{(0)}$) admits a unique ground state and first-excited state will continue to do so for all $\boldsymbol{h} \in \mathcal{F}_H$ as long as H is small enough that the additional external field does not alter the low-energy spectrum, i.e., as long as $H < \frac{1}{2}\min\{\Delta E_{01}, \Delta E_{12}\}$ where ΔE_{01} and ΔE_{12} denote the initial energy gaps between the ground and first-excited states, and first-excited and second-excited states, respectively.

In the case described above, we find that the low-temperature magnetization is maximized by widening or narrowing the energy gap between $\boldsymbol{\sigma}^0$ and $\boldsymbol{\sigma}^1$. Under an external field $\boldsymbol{h}^{(0)} + \boldsymbol{h}$, for $\boldsymbol{h} \in \mathcal{F}_H$, the energy gap between the ground and first-excited states takes the form:

$$\begin{aligned}\Delta E(\boldsymbol{h}) &= E_{\boldsymbol{h}}(\boldsymbol{\sigma}^1) - E_{\boldsymbol{h}}(\boldsymbol{\sigma}^0) \\ &= \Delta E_{01} + (\boldsymbol{\sigma}^0 - \boldsymbol{\sigma}^1)^T \boldsymbol{h}. \end{aligned} \tag{13}$$

For sufficiently low temperatures, we can write the total magnetization of the system in terms of $\Delta E(\boldsymbol{h})$:

$$\begin{aligned} M(\boldsymbol{h}^{(0)} + \boldsymbol{h}) &\approx \frac{M_0 + M_1 e^{-\frac{1}{T}\Delta E(\boldsymbol{h})}}{1 + e^{-\frac{1}{T}\Delta E(\boldsymbol{h})}} \\ &\approx M_0 + (M_1 - M_0)\, e^{-\frac{1}{T}\Delta E(\boldsymbol{h})}, \end{aligned} \tag{14}$$

where $M_0 = \sum_i \sigma_i^0$ and $M_1 = \sum_i \sigma_i^1$ denote the magnetizations of the ground and first-excited states, respectively. Thus, the low-temperature optimal external fields, i.e., the external fields which maximize the low-temperature magnetization, take the form:

$$\begin{aligned} \boldsymbol{h}^* &= \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} (M_1 - M_0)\, e^{-\frac{1}{T}\Delta E(\boldsymbol{h})} \\ &\equiv \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} -\operatorname{sign}[M_1 - M_0]\, \Delta E(\boldsymbol{h}) \\ &\equiv \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} \operatorname{sign}[M_0 - M_1]\, (\boldsymbol{\sigma}^0 - \boldsymbol{\sigma}^1)^T \boldsymbol{h}. \end{aligned} \tag{15}$$

We first note that Eq. (15) is linear in $\boldsymbol{h}$ and is hence exactly solvable by efficient and scalable linear programming techniques. Thus, in the case of a unique ground state and a unique first-excited state, the efficiency of calculating the low-T solutions to IIM is limited only by the computational complexity of identifying the ground and first-excited states. While calculating the ground state for a general Ising system is known to be NP-hard [2], we note that there are efficient algorithms based on graph cuts for approximately minimizing the energy of general Ising systems within a known factor [5]. Furthermore, for ferromagnetic systems in arbitrary external fields, there are graph cut algorithms for provably and efficiently calculating the low-energy states [20]. Since such graph-cut algorithms are primarily used in computer vision, Eq. (15) presents an interesting connection between IIM and computer vision problems in machine learning.

In addition to its interesting algorithmic properties, Eq. (15) also reveals a major insight into the structure of low-T optimal external fields. Depending on the sign of $(M_0 - M_1)$, the optimal external field $\boldsymbol{h}^*$ either maximizes the energy gap (in the case $M_0 > M_1$) or minimizes the energy gap (in the case $M_0 < M_1$). Furthermore, since the change in the energy gap equals $(\boldsymbol{\sigma}^0 - \boldsymbol{\sigma}^1)^T \boldsymbol{h}$, it is clear that to maximize or minimize the energy gap, external field resources should be focused on nodes i for which $\sigma_i^0 \neq \sigma_i^1$. Thus, as long as $M_0 \neq M_1$, the optimal external fields focus H on the nodes with opposite parity between the ground and first-excited states. The result is summarized in Theorem 3.

Theorem 3. Consider an Ising system, described by J and $\boldsymbol{h}^{(0)}$, for which there exist unique ground and first-excited states $\boldsymbol{\sigma}^0$ and $\boldsymbol{\sigma}^1$ for all feasible external fields $\boldsymbol{h} \in \mathcal{F}_H$, for some external field budget H. Letting $N_{LC} = \{i \in N : \sigma_i^0 \neq \sigma_i^1\}$ denote the set of loosely-connected nodes with opposite parity between the ground and first-excited states, if $M_0 \neq M_1$, then for sufficiently low temperatures, all optimal external fields $\boldsymbol{h}^* \in \mathcal{F}_H$ satisfy $\sum_{i\in N_{LC}} |h_i^*| = H$, and hence $h_i^* = 0$ for all $i \notin N_{LC}$.

Here we emphasize the structural dissimilarity between the high- and low-T optimal external fields. For general Ising systems, at high temperatures, we found in Section III that the optimal external fields focus on the nodes of highest degree in the network. We contrast this picture with the low-temperature limit, where, if the system admits unique ground and first-excited states for all $\boldsymbol{h} \in \mathcal{F}_H$, the optimal external fields focus on the loosely-connected nodes, which have opposite parity between the ground and first-excited states. In ferromagnetic systems, if the network is well-connected in the sense that the energy associated with flipping a set of nodes from the ground state tends to increase with the size of the set, then the first-excited state will differ from the ground state in only one or two entries, corresponding to nodes of low degree in the network. In Section VI, we verify this intuition, showing on a number of random and real-world networks that the low-T optimal external fields focus on nodes of low degree. Finally, we point the reader to Appendix B for an analytic description of the low-T optimal

external fields for general Ising systems and general external field budgets.

IV.2. A low-T analytic solution under the mean-field approximation

In order to connect the above results to previous work on the IIM problem, in this subsection we compare the optimal external fields to the mean-field optimal external fields in the low-temperature limit. As argued in the previous subsection, formulating a general description of the low-temperature mean-field magnetization is unilluminating. Thus, to ease the discussion, we limit our study to the mean-field analogue of the assumptions made in the previous subsection; namely, when the system admits a unique ground state and a unique node with the smallest effective field for all external fields $\boldsymbol{h} \in \mathcal{F}_H$.

To clarify our assumptions, consider an Ising system described by J, $\boldsymbol{h}^{(0)}$, and T and an external field budget H. Our first assumption is that there exists a unique ground state $\boldsymbol{\sigma}^0$ for all $\boldsymbol{h} \in \mathcal{F}_H$. For low temperatures, under an external field $\boldsymbol{h}^{(0)} + \boldsymbol{h}$, for $\boldsymbol{h} \in \mathcal{F}_H$, the mean-field magnetization of node i takes the form:

$$\begin{aligned} m_i &\approx \tanh\left[\frac{1}{T}\left(\sum_j J_{ij}\sigma_j^0 + h_i^{(0)} + h_i\right)\right] \\ &\approx \sigma_i^0\left(1 - 2e^{-\frac{1}{T}2\sigma_i^0\left(\sum_j J_{ij}\sigma_j^0 + h_i^{(0)} + h_i\right)}\right) \\ &\equiv \sigma_i^0\left(1 - 2e^{-\frac{1}{T}2\tilde{h}_i}\right), \end{aligned} \tag{16}$$

where $\tilde{h}_i = \sigma_i^0\left(\sum_j J_{ij}\sigma_j^0 + h_i^{(0)} + h_i\right)$ is the *effective field* felt by node i, and the second approximation follows by Taylor expansion. Since $\sigma_i^0 = \text{sign}\left[\sum_j J_{ij}\sigma_j^0 + h_i^{(0)} + h_i\right]$, we note that $\tilde{h}_i \geq 0$ for all $i \in N$. Furthermore, we note that the nodes i for which $|m_i - \sigma_i^0|$ is greatest are the nodes with the smallest effective fields $\tilde{h}_i$. Thus, our second and final assumption is that there exists a unique node $i^* = \arg\min_i \tilde{h}_i$ for all $\boldsymbol{h} \in \mathcal{F}_H$. As suggested by the notation, we will show that the low-T mean-field magnetization is maximized by focusing on node i^*.

Under the above assumptions, and for an external field $\boldsymbol{h}^{(0)} + \boldsymbol{h}$, where $\boldsymbol{h} \in \mathcal{F}_H$, the low-temperature mean-field magnetization takes the form:

$$\begin{aligned} M_{MF}(\boldsymbol{h}^{(0)} + \boldsymbol{h}) &\approx \sum_i \sigma_i^0\left(1 - 2e^{-\frac{1}{T}2\tilde{h}_i}\right) \\ &\approx M_0 - 2\sigma_{i^*}^0 e^{-\frac{1}{T}2\tilde{h}_{i^*}}, \end{aligned} \tag{17}$$

since all other terms will be exponentially suppressed in the $T \to 0$ limit. Thus, for an external field budget H, the low-temperature optimal external field under the mean-field approximation is given by:

$$\begin{aligned} \boldsymbol{h}_{MF}^* &= \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} -2\sigma_{i^*}^0 e^{-\frac{1}{T}2\tilde{h}_{i^*}} \\ &\equiv \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} \sigma_{i^*}^0 \tilde{h}_{i^*} \\ &= \arg\max_{\boldsymbol{h}\in\mathcal{F}_H} \sum_j J_{i^*j}\sigma_j^0 + h_{i^*}^{(0)} + h_{i^*} \\ &= H\boldsymbol{e}_{i^*}, \end{aligned} \tag{18}$$

where $\boldsymbol{e}_i \in \mathbb{R}^n$ denotes the vector of all zeros, but for a one in the i^{th} conmponent. Thus, when there exists a unique ground state and a unique node $i^* = \arg\min_i \tilde{h}_i$ for all $\boldsymbol{h} \in \mathcal{F}_H$, the low-temperature mean-field magnetization is maximized by focusing H on node i^*.

To illustrate the differences between $\boldsymbol{h}^*$ and $\boldsymbol{h}_{MF}^*$ at low temperatures, we consider the small ferromagnetic system shown in Figure 3, where thin edges represent bonds of weight 1, thick edges represent bonds of weight 3, and the system feels an initial external field $\boldsymbol{h}^{(0)} = (0, 0, 5, 0)^T$. For $H < \frac{1}{2}$, the unique ground state is all-up $\boldsymbol{\sigma}^0 = (1, 1, 1, 1)^T$, and the unique first-excited state is $\boldsymbol{\sigma}^1 = (-1, -1, 1, 1)^T$ for all $\boldsymbol{h} \in \mathcal{F}_H$. The low-$T$ optimal external fields $\boldsymbol{h}^*$ focus on the loosely-connected nodes with opposite parity between the ground and first-excited states (nodes 1 and 2). On the other hand, the mean-field low-T optimal external field $\boldsymbol{h}_{MF}^*$ focuses on the node with the smallest effective field (node 4). We also note for comparison that the high-T optimal external field focuses on the node of highest degree (node 3), independent of $\boldsymbol{h}^{(0)}$ and H.

V. A GRADIENT ASCENT ALGORITHM

In addition to analytically exploring the structure optimal external fields in the high- and low-T limits, we are also interested in developing efficient algorithmic tools for studying the structure of optimal external fields at all temperatures. In this section we present a novel algorithm that efficiently calculates a local maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ in $\mathcal{F}_H$ for general Ising systems at arbitrary temperatures and for general external field budgets. For concreteness, by efficient we mean that, for any accuracy parameter $\epsilon > 0$, our algorithm finds an ϵ-approximation to a local maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ in $O(1/\epsilon)$ iterations. In the special case of a ferromagnetic network ($J \geq 0$) in a nonnegative initial external field ($\boldsymbol{h}^{(0)} \geq 0$), we show that the magnetization is non-decreasing and concave in $\boldsymbol{h}$, and hence that our algorithm is guaranteed to converge to a global maximum of the magnetization, efficiently solving the IIM problem.

While previous work has focused on maximizing the mean-field magnetization [24], our algorithm is the first to provide performance guarantees with respect to the exact magnetization. Our algorithm, summarized in Algorithm 1, is based on standard projected gradient ascent techniques. The algorithm is initialized at some feasible external field $\boldsymbol{h}^{(1)} \in \mathcal{F}_H$, and subsequently

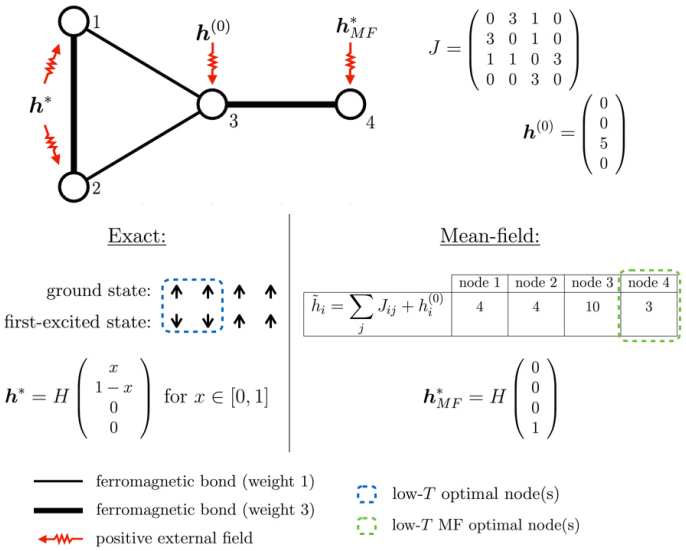
1
2
3
4
$\boldsymbol{h}^{(0)}$
$\boldsymbol{h}^{*}_{MF}$
$\boldsymbol{h}^{*}$
$J = \begin{pmatrix} 0 & 3 & 1 & 0 \\ 3 & 0 & 1 & 0 \\ 1 & 1 & 0 & 3 \\ 0 & 0 & 3 & 0 \end{pmatrix}$
$\boldsymbol{h}^{(0)} = \begin{pmatrix} 0 \\ 0 \\ 5 \\ 0 \end{pmatrix}$
Exact:
ground state:
first-excited state:
$\boldsymbol{h}^{*} = H \begin{pmatrix} x \\ 1-x \\ 0 \\ 0 \end{pmatrix}$ for $x \in [0,1]$
Mean-field:
node 1 node 2 node 3 node 4
$\tilde{h}_i = \sum_j J_{ij} + h_i^{(0)}$ 4 4 10 3
$\boldsymbol{h}^{*}_{MF} = H \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$
ferromagnetic bond (weight 1)
ferromagnetic bond (weight 3)
positive external field
low-T optimal node(s)
low-T MF optimal node(s)

FIG. 3. A small ferromagnetic network with bonds of weight 1 and 3 represented by thin and thick edges, respectively, under a positive initial external field $\boldsymbol{h}^{(0)} = (0, 0, 5, 0)^T$. For $H < \frac{1}{2}$, the unique ground state is all-up $\boldsymbol{\sigma}^0 = (1, 1, 1, 1)^T$ and the unique first-excited state is $\boldsymbol{\sigma}^1 = (-1, -1, 1, 1)^T$ for all $\boldsymbol{h} \in \mathcal{F}_H$. The low-$T$ optimal external fields apply H to the loosely-connected nodes (nodes 1 and 2), while the mean-field optimal external field focuses H on the node with the smallest effective field (node 4).

Algorithm 1: Projected gradient ascent

Input: An Ising system described by J, $\boldsymbol{h}^{(0)}$, and T; an external field budget H; and an accuracy parameter $\epsilon > 0$

Output: An external field $\boldsymbol{h}$ that approximates a local maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ in $\mathcal{F}_H$

$t = 1$; $\boldsymbol{h}^{(1)} \in \mathcal{F}_H$; $\alpha \in (0, \frac{1}{L})$;

repeat

$\frac{\partial M}{\partial h_i} = \sum_j \chi_{ij}$ for all $i \in N$ under $\boldsymbol{h}^{(0)} + \boldsymbol{h}^{(t)}$;

$\boldsymbol{h}^{(t+1)} = P_{\mathcal{F}_H}\left[\boldsymbol{h}^{(t)} + \alpha \nabla_{\boldsymbol{h}} M(\boldsymbol{h}^{(0)} + \boldsymbol{h}^{(t)})\right]$;

t++;

until $\left|M(\boldsymbol{h}^{(0)} + \boldsymbol{h}^{(t+1)}) - M(\boldsymbol{h}^{(0)} + \boldsymbol{h}^{(t)})\right| \leq \epsilon$;

$\boldsymbol{h} = \boldsymbol{h}^{(t+1)}$;

steps along the direction of the gradient of the magnetization, projected onto $\mathcal{F}_H$, until a local maximum is reached. In particular, at each iteration t, we calculate the gradient of the magnetization, $\frac{\partial M}{\partial h_i} = \sum_j \chi_{ij} = \frac{1}{T}\sum_j \left(\langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle\right)$, under the external field $\boldsymbol{h}^{(0)} + \boldsymbol{h}^{(t)}$ and project the gradient onto the space of feasible external fields. We note that the projection operator, denoted $P_{\mathcal{F}_H}$, is well-defined since $\mathcal{F}_H$ is convex. Stepping along the direction of the projected gradient with step size $\alpha \in (0, \frac{1}{L})$, where L is a Lipschitz constant of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ (which is well-defined since M is smooth for finite systems), Algorithm 1 is known to converge to an ϵ-approximation to a local maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ in $O(1/\epsilon)$ iterations [41].

We now show, in the special case of a ferromagnetic network ($J \geq 0$) in a nonnegative external field ($\boldsymbol{h}^{(0)} \geq 0$), that the magnetization $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ is non-decreasing and concave in $\boldsymbol{h}$, for $\boldsymbol{h} \geq 0$, and hence is efficiently maximized by Algorithm 1. The main result is summarized in Theorem 4.

Theorem 4. Consider a ferromagnetic network ($J \geq 0$) in a nonnegative initial external field ($\boldsymbol{h}^{(0)} \geq 0$) with arbitrary temperature $T > 0$ and arbitrary external field budget $H > 0$. Then, choosing $\boldsymbol{h}^{(1)} \geq 0$, Algorithm 1 converges to an ϵ-approximation to IIM in $O(1/\epsilon)$ iterations for any accuracy parameter $\epsilon > 0$.

To establish Theorem 4, we first note that the Fortuin-Kasteleyn-Ginibre inequality [10] states that, for any ferromagnetic system, regardless of external field, $\chi_{ij} = \frac{\partial m_i}{\partial h_j} \geq 0$ for all $i, j \in N$, and hence $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ is non-decreasing in $\boldsymbol{h}$. This implies two important facts: (i) there exists a global maximum in $\mathcal{F}_H$ for which $\boldsymbol{h} \geq 0$, and (ii) if Algorithm 1 is initialized such that $\boldsymbol{h}^{(1)} \geq 0$, then we will have $\boldsymbol{h}^{(t)} \geq 0$ for all subsequent t. Secondly, the Griffiths-Hurst-Sherman inequality [16] guarantees that, for any ferromagnetic system in a nonnegative external field, $\frac{\partial^2 m_i}{\partial h_j \partial h_k} \leq 0$ for all $i, j, k \in N$, and hence $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ is concave in $\boldsymbol{h}$ for $\boldsymbol{h} \geq 0$. This implies that any local maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$ for $\boldsymbol{h} \geq 0$ is a global maximum for $\boldsymbol{h} \geq 0$ and hence, by fact (i) above, is a global maximum among all $\boldsymbol{h} \in \mathcal{F}_H$. Thus, in the case of a ferromagnetic system in a nonnegative initial external field, choosing $\boldsymbol{h}^{(1)} \geq 0$, Algorithm 1 will converge to a global maximum of $M(\boldsymbol{h}^{(0)} + \boldsymbol{h})$, efficiently solving the IIM problem. In the following section we apply Algorithm 1 on a number of random and real-world networks to study the structure of optimal external fields across a range of temperatures.

Remark. While Algorithm 1 is efficient in that it converges to a local maximum in $O(1/\epsilon)$ iterations, the total run-time is limited by the calculation of χ and M at each step. Since exact Ising calculations involve sums over $\{\pm 1\}^n$, which is exponential in the size of the system, an exact implementation of Algorithm 1 is limited to relatively small networks. To overcome this exponential dependence on system size, in the next section we use Monte Carlo simulations to approximate χ and M at each iteration of Algorithm 1.

VI. NUMERICAL EXPERIMENTS

In this section, we present numerical experiments to probe the structure of optimal external fields and to compare the performance of Algorithm 1 with the state-of-the-art mean-field-based algorithm in [24]. For a number of random and real-world networks, we verify our analytic results from Sections III and IV, showing that the optimal external fields undergo phase shifts from focusing on high-degree nodes at high temperatures to focusing on

low-degree nodes at low temperatures.

Since an exact implementation of Algorithm 1 is limited to relatively small networks, to study large systems we employ an approximation to Algorithm 1 in which χ and M are calculated via Monte Carlo simulations at each iteration. For simplicity, we refer to the exact implementation of Algorithm 1 as Algorithm 1(exact) and the Monte Carlo approximation as Algorithm 1(MC). On small random networks, we confirm that Algorithm 1(exact) outperforms both Algorithm 1(MC) as well as the mean-field-based algorithm in [24], as expected. We also apply Algorithm 1(MC) on a number of large random and real-world networks, exhibiting its scalability and comparing the performance of Algorithm 1(MC) with the mean-field-based algorithm in [24].

VI.1. Small networks

We begin by studying the structure of optimal external fields, calculated exactly via Algorithm 1(exact), and comparing them with approximately optimal external fields, calculated via Monte Carlo simulations (Algorithm 1(MC)) and the mean-field approximation [24]. Since the algorithms we consider only have performance guarantees for ferromagnetic systems in nonnegative external fields, we limit our study to simple networks ($J_{ij} \in \{0, 1\}$) with no initial external field ($\boldsymbol{h}^{(0)} = 0$), and, for simplicity, we set $H = 1$.

We first consider an Erdös-Rényi random network consisting of n=15 nodes, where each node is connected to each other node with probability 0.3. The resulting system, shown in Figure 4, has maximum degree 8, minimum degree 2, and mean-field critical temperature $T_{c,MF} = \rho(J) = 5.45$ (see [24]), which sets the temperature scale of the system. For finite systems, the exact magnetization does not exhibit a phase transition; however, under a non-zero external field, any finite ferromagnetic system will exhibit a phase shift (a smooth transformation) marked by a peak in the susceptibility. We denote the temperature corresponding to maximum susceptibility by $\tilde{T}_c$, and for our Erodös-Rényi network with $H = 1$ we have $\tilde{T}_c = 1.26$.

In Figure 4(a), we compare the optimal magnetization, denoted $M(\boldsymbol{h}^*)$, with the magnetizations under external fields $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform} = \frac{H}{n}(1, \ldots, 1)^T$, where $\boldsymbol{h}^*$ is calculated using Algorithm 1(exact), $\boldsymbol{h}^*_{MC}$ is calculated using Algorithm 1(MC) with Glauber dynamics consisting of 10,000 rounds of burn-in and 10,000 rounds of averaging at each iteration [22], and $\boldsymbol{h}^*_{MF}$ maximizes the mean-field magnetization [24]. We see clearly that $\boldsymbol{h}^*$ provides the largest magnetization, as expected, with $\boldsymbol{h}^*_{MC}$ admitting performance losses of up to 10% and $\boldsymbol{h}^*_{MF}$ suffering performance losses of up to 35% at intermediate temperatures. Furthermore, we find that Algorithm 1(MC) performs worst near $\tilde{T}_c$, likely due to critical slowing down of the Monte Carlo dynamics [31]. We also note that the external fields are all near-optimal at high and low temperatures, since, for all $\boldsymbol{h} > 0$, the magnetization approaches 0 as $T \to \infty$ and the magnetization approaches n as $T \to 0$.

Figure 4(b) compares the phase shift in $\boldsymbol{h}^*$ with the phase shifts in $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ by calculating the average degrees of the external fields, defined by $\frac{\sum_i d_i h_i}{H}$, over a range of temperatures. We find that $\boldsymbol{h}^*$, $\boldsymbol{h}^*_{MC}$, and $\boldsymbol{h}^*_{MF}$ all exhibit phase shifts from focusing on high-degree nodes at high temperatures to focusing on low-degree nodes at low temperatures. However, each external field shifts at a different temperature, with $\boldsymbol{h}^*$ making the transition at the lowest temperature, followed by $\boldsymbol{h}^*_{MC}$ and then $\boldsymbol{h}^*_{MF}$ near $T_{c,MF}$.

We now consider a small preferential attachment network (also known as a Barabási-Albert network) consisting of $n = 15$ nodes, shown in Figure 5. Preferential attachment networks exhibit power-law degree distributions widely observed in social networks [1]. Our system has a hub of maximum degree 10, five nodes of minimum degree 2, phase shift temperature $\tilde{T}_c = 1.18$, and mean-field critical temperature $T_{c,MF} = 4.51$. In Figure 5(a), we compare the magnetizations under external fields $\boldsymbol{h}^*$, $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$ over a range of temperatures. As in the Erodös-Rényi network, $\boldsymbol{h}^*$ achieves the highest magnetization for all temperatures, while $\boldsymbol{h}^*_{MC}$ performs second-best with performance losses limited to 4% and localized near $\tilde{T}_c$, with $\boldsymbol{h}^*_{MF}$ suffering performance losses of up to 40%. In Figure 5(b), we study the average degrees of $\boldsymbol{h}^*$, $\boldsymbol{h}^*_{MC}$, and $\boldsymbol{h}^*_{MF}$ over a range of temperatures, finding that all three exhibit phase shifts from focusing on the hub node at high temperatures to spreading H among the minimum degree nodes at low temperatures. Furthermore, as in the Erodös-Rényi network, we find that $\boldsymbol{h}^*$ transitions at the lowest temperature, followed by $\boldsymbol{h}^*_{MC}$ and then $\boldsymbol{h}^*_{MF}$ near $T_{c,MF}$.

VI.2. Large networks

We now shift our focus to large networks in order to exhibit the scalability of Algorithm 1(MC) and to study the structure of optimal external fields in large-scale systems. However, in scaling to large systems we lose two important tools: (i) we lose the use of Algorithm 1(exact) to exactly calculate optimal external fields, and (ii) we lose the ability to exactly calculate magnetizations and hence evaluate performance. Thus, in order to compare the performance of Algorithm 1(MC) with the mean-field-based algorithm in [24], we calculate the magnetizations using 20 Glauber Monte Carlo runs with 5,000 rounds of burn-in and 10,000 rounds of averaging each. As in the previous subsection, we restrict our attention to simple networks ($J_{ij} \in \{\pm 1\}$) with no initial external field ($\boldsymbol{h}^{(0)} = 0$).

We first consider an Erodös-Rényi network, shown in Figure 6, consisting of $n = 200$ nodes with edge probability 0.04. The resulting network has one node of maximum degree 17, one node of minimum degree 2, and mean-field critical temperature $T_{c,MF} = 9.15$. In Figure 6(a), we compare of the magnetizations due to $\boldsymbol{h}^*_{MC}$,

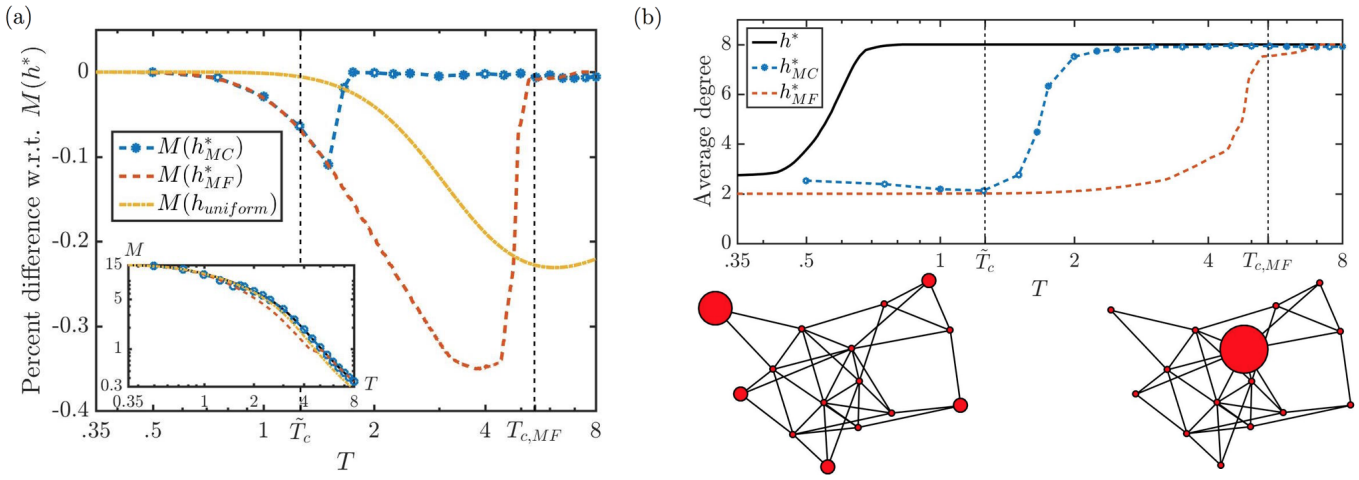

FIG. 4. We consider an Erdös-Rényi random network composed of 15 nodes with edge probability 0.3 and couplings of weight 1, under initial external field $\boldsymbol{h}^{(0)} = 0$. (a) For $H = 1$, we compare $M(\boldsymbol{h}^*)$ with the magnetizations under $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$, confirming that $\boldsymbol{h}^*$ achieves the highest magnetization for all temperatures. (b) We verify that $\boldsymbol{h}^*$, $\boldsymbol{h}^*_{MC}$, and $\boldsymbol{h}^*_{MF}$ all undergo phase shifts from focusing on high- to low-degree nodes as the temperature decreases.

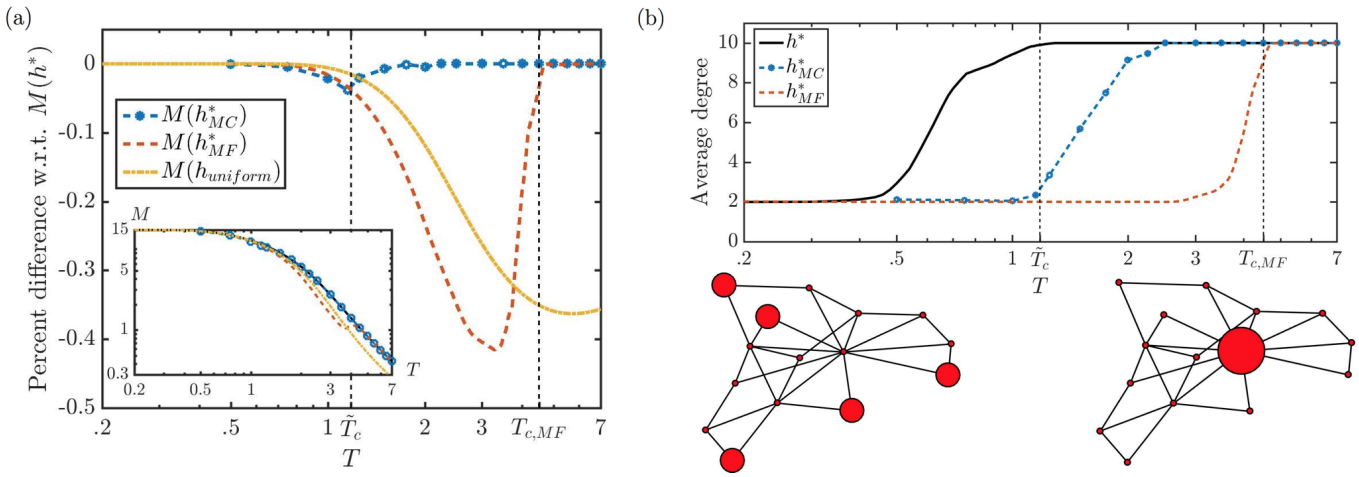

FIG. 5. We consider a preferential attachment network composed of 15 nodes and couplings of weight 1, under initial external field $\boldsymbol{h}^{(0)} = 0$. (a) For $H = 1$, we compare $M(\boldsymbol{h}^*)$ with the magnetizations under $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$, verifying that $\boldsymbol{h}^*$ achieves the highest magnetization for all temperatures. (b) We verify that $\boldsymbol{h}^*$, $\boldsymbol{h}^*_{MC}$, and $\boldsymbol{h}^*_{MF}$ all undergo phase shifts from focusing on high- to low-degree nodes as the temperature decreases.

$\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$ for an external field budget $H = 10$, where the error bars represent one standard deviation from the mean for each batch of 20 Monte Carlo simulations. Across most temperatures we find that $\boldsymbol{h}^*_{MC}$ performs favorably compared with $\boldsymbol{h}^*_{MF}$, especially near $T_{c,MF}$, where $\boldsymbol{h}^*_{MF}$ admits performance losses of up to 40%.

In Figure 6(b), we investigate the structure of $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ across a range of temperatures, specifically studying the average degree and the participation ratio, defined by $\frac{(\sum_i h_i)^2}{\sum_i h_i^2}$, which measures the effective number of nodes in the network over which H is spread. As illustrated in the small networks, we find that both $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ exhibit phase shifts from focusing on high-degree to low-degree nodes as the temperature decreases, with $\boldsymbol{h}^*_{MF}$ transitioning near $T_{c,MF}$ and $\boldsymbol{h}^*_{MC}$ transitioning at a lower temperature. We also find for both high and low temperatures that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ focus on only one or two nodes, while near the phase shifts they spread H among many nodes. This indicates that the optimal external fields at intermediate temperatures involve many nodes not included in either the high-T or low-T optimal external fields.

Next, we consider a large preferential attachment network, shown in Figure 7, consisting of $n = 200$ nodes with one node of maximum degree 47, 99 nodes of minimum degree 2, and mean-field critical temperature $T_{c,MF} = 8.67$. In Figure 7(a), we compare the magnetization of the system under external fields $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$ for an external field budget $H = 10$, finding that Algorithm 1(MC) compares favorably with the mean-field-based algorithm in [24] across most tempera-

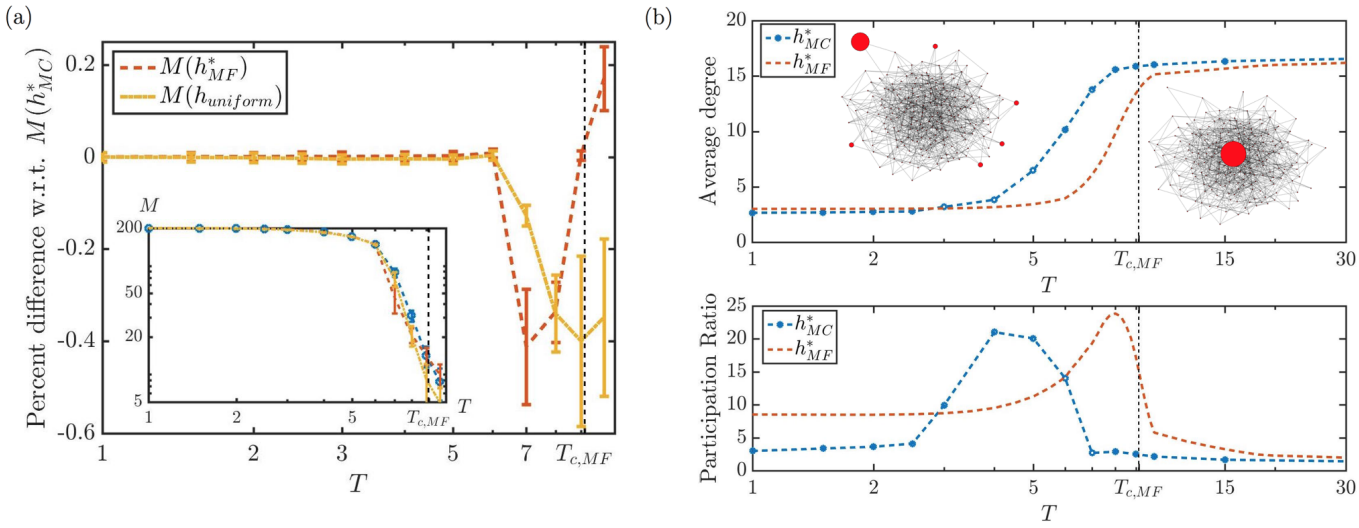

FIG. 6. We consider an Erdös-Rényi random network composed of 200 nodes with edge probability 0.04 and couplings of weight 1, under initial external field $\boldsymbol{h}^{(0)} = 0$. (a) For $H = 10$, we compare the magnetizations under $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$, finding that $\boldsymbol{h}^*_{MC}$ achieves the highest magnetization for most temperatures. (b) We verify that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ both undergo phase shifts from focusing on high- to low-degree nodes as the temperature decreases. Furthermore, we find for high and low temperatures that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ only focus on a few nodes, while near the phase shifts they spread H among many nodes.

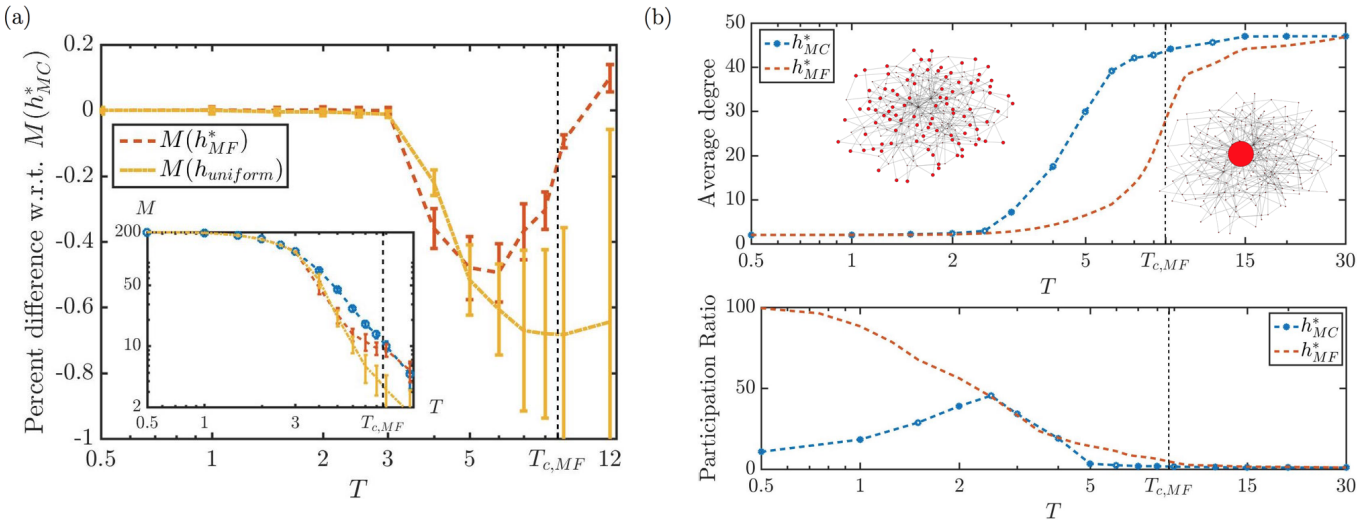

FIG. 7. We consider a preferential attachment network composed of 200 nodes and couplings of weight 1, under initial external field $\boldsymbol{h}^{(0)} = 0$. (a) For $H = 10$, we compare the magnetizations under $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$, finding that $\boldsymbol{h}^*_{MC}$ achieves the highest magnetization for most temperatures. (b) We verify that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ both undergo phase shifts from focusing on high- to low-degree nodes as the temperature decreases. Furthermore, we find for high temperatures that both $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ focus on the node of highest degree, while at intermediate temperatures, $\boldsymbol{h}^*_{MC}$ spreads H among a number of nodes. At low temperatures, $\boldsymbol{h}^*_{MF}$ spreads H among the 99 nodes of lowest degree.

tures. In Figure 7(b), we study the structure of $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ in the preferential attachment network, confirming that both external fields exhibit phase shifts from high- to low-degree nodes as the temperature decreases. Furthermore, looking at the participation ratio, we find that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ both focus on the single node of highest degree at high temperatures, and, as in the Erodös-Rényi network, $\boldsymbol{h}^*_{MC}$ spreads H among many nodes near its phase shift. However, in contrast to the Erodös-Rényi network, $\boldsymbol{h}^*_{MF}$ continues to spread H among an increasing number of nodes as the temperature decreases, eventually spreading H among the 99 nodes of lowest degree.

Finally, we consider a real-world collaboration network, shown in Figure 8, consisting of $n = 904$ high-energy physicists and 2345 undirected, unweighted edges, where each edge represents the co-authorship of a paper posted to the arXiv [21]. We note that co-authorship networks are known to capture many of the key structural features of social networks [30], and our network has a maximum degree of 38, 127 nodes of minimum degree 1,

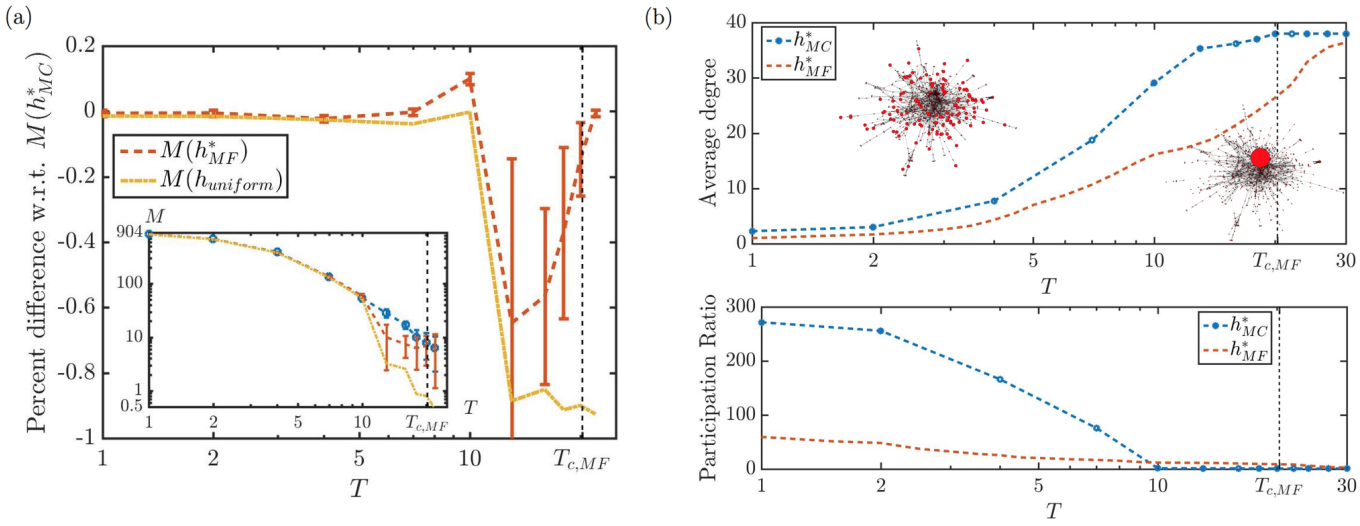

FIG. 8. We consider a collaboration network consisting of 904 physicists where the edges represent co-authorships on papers posted to the arXiv. (a) For $H = 20$, we compare the magnetizations under $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$, finding that $\boldsymbol{h}^*_{MC}$ performs well compared with $\boldsymbol{h}^*_{MF}$. (b) We confirm that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ both undergo phase shifts from focusing on high-degree nodes to spreading H among low-degree nodes as the temperature decreases.

and a mean-field critical temperature $T_{c,MF} = 20.26$. In Figure 8(a), we compare the magnetization of the social network under external fields $\boldsymbol{h}^*_{MC}$, $\boldsymbol{h}^*_{MF}$, and $\boldsymbol{h}_{uniform}$ for an external field budget $H = 20$, finding that Algorithm 1(MC) performs well compared with the mean-field based algorithm in [24]. We also note that the error bars for the percent difference of $M(\boldsymbol{h}_{uniform})$ with respect to $M(\boldsymbol{h}^*_{MC})$ have been intentionally omitted to improve the presentation. In Figure 8(b), we confirm that $\boldsymbol{h}^*_{MC}$ and $\boldsymbol{h}^*_{MF}$ undergo phase shifts from focusing on the node of highest degree at high temperatures to spreading H among the nodes of lowest degree at low temperatures. Interestingly, at low temperatures we find that $\boldsymbol{h}_{MC}$ spreads H among many more nodes than $\boldsymbol{h}_{MF}$.

VII. CONCLUSION

In this paper we explore the Ising influence maximization problem, which has recently been proposed as a framework for studying the control and optimization of complex, reverberant opinion dynamics in large social networks. Traditionally, influence maximization in social networks has been studied in the context of contagion models and irreversible processes, so the IIM problem, which has a natural physical interpretation as the maximization of the magnetization given a budget of external magnetic field, marks an important departure from the existing paradigm and bridges a gap between the IM and Ising literatures. We are the first to consider the IIM problem in general Ising systems with negative couplings and negative external fields, accounting for anti-correlated opinions and competitive marketing, respectively, two important features found in real-world social networks.

In this paper we focus on describing the structure of optimal external fields (i.e. external fields which maximize the magnetization), and we provide novel analytic and algorithmic tools for identifying influential nodes in general complex Ising systems. Remarkably, we show analytically that the optimal external fields exhibit phase shifts from intuitively focusing on high-degree nodes at high temperatures to counterintuitively focusing on "loosely-connected" nodes, which are weakly energetically bound to the ground state, at low temperatures. In order to connect our results with previous work that has focused on maximizing the mean-field magnetization, we also provide novel analytic descriptions of the mean-field optimal external fields, showing that they too exhibit a phase shift from focusing on high-degree nodes at high temperatures to focusing on the nodes with smallest effective fields at low temperatures.

We are also the first to study the Ising influence maximization problem directly (without mean-field approximations), and we provide a novel and efficient algorithm for finding local maxima of the magnetization in general Ising systems. Furthermore, for ferromagnetic systems in non-negative initial external fields, we prove that our algorithm converges to a global maximum, and hence efficiently solves the IIM problem. Finally, we apply our algorithm on large random and real-world networks, verifying the existence of phase shifts in the optimal external fields and comparing the performance of our algorithm with the state-of-the-art mean-field-based algorithm in [24].

We believe that this paper opens the door for many exciting directions for future work. Many of the analytic and algorithmic tools developed to study Ising networks, ranging from renormalization group methods to advanced mean-field techniques, can now be applied in a wide array of contexts to study influence maximization in social networks. Furthermore, we believe that the

IIM problem opens a new chapter for the Ising model by introducing a framework for exploring the control and optimization of complex Ising systems.

ACKNOWLEDGEMENTS

We acknowledge support from the U.S. National Science Foundation, the Air Force Office of Scientific Research, and the Department of Transportation.

APPENDIX A: HIGH-T SUSCEPTIBILITY

In what follows, we derive the first- and second-order terms in the high-temperature expansion of the susceptibility $\frac{1}{\beta}\chi_{ij} = \langle\sigma_i\sigma_j\rangle - \langle\sigma_i\rangle\langle\sigma_j\rangle$ in Eq. (4). As noted in the main text, for any function of the spins $f(\boldsymbol{\sigma})$ we have

$$\langle f(\boldsymbol{\sigma})\rangle_{\beta=0} = \frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}} f(\boldsymbol{\sigma}), \tag{19}$$

where the sum runs over all spin configurations. We also note that

$$\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}} \sigma_i\sigma_j = \delta_{ij}, \tag{20}$$

$$\begin{aligned}\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}} \sigma_i\sigma_j\sigma_k\sigma_\ell &= \delta_{ij}\delta_{k\ell} + \delta_{ik}\delta_{j\ell} + \delta_{i\ell}\delta_{jk} - 2\delta_{ijk\ell}\\ &:= C_{ijk\ell},\end{aligned} \tag{21}$$

$$\begin{aligned}\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}} \sigma_i\sigma_j\sigma_k\sigma_\ell\sigma_m\sigma_p &= \delta_{ij}(\delta_{k\ell}\delta_{mp} + \delta_{km}\delta_{\ell p} + \delta_{kp}\delta_{\ell m})\\ &+ \delta_{ik}(\delta_{j\ell}\delta_{mp} + \delta_{jm}\delta_{\ell p} + \delta_{jp}\delta_{\ell m})\\ &+ \delta_{i\ell}(\delta_{jk}\delta_{mp} + \delta_{jm}\delta_{kp} + \delta_{jp}\delta_{km})\\ &+ \delta_{im}(\delta_{jk}\delta_{\ell p} + \delta_{j\ell}\delta_{kp} + \delta_{jp}\delta_{k\ell})\\ &+ \delta_{ip}(\delta_{jk}\delta_{\ell m} + \delta_{j\ell}\delta_{km} + \delta_{jm}\delta_{k\ell})\\ &- 2\left[\delta_{ij}\delta_{k\ell mp} + \delta_{ik}\delta_{j\ell mp} + \delta_{i\ell}\delta_{jkmp}\right.\\ &+ \delta_{im}\delta_{jk\ell p} + \delta_{ip}\delta_{jk\ell m} + \delta_{jk}\delta_{i\ell mp}\\ &+ \delta_{j\ell}\delta_{ikmp} + \delta_{jm}\delta_{ik\ell p} + \delta_{jp}\delta_{ik\ell m}\\ &+ \delta_{k\ell}\delta_{ijmp} + \delta_{km}\delta_{ij\ell p} + \delta_{kp}\delta_{ij\ell m}\\ &\left.+\delta_{\ell m}\delta_{ijkp} + \delta_{\ell p}\delta_{ijkm} + \delta_{mp}\delta_{ijk\ell}\right]\\ &+ 16\delta_{ijk\ell mp}\\ &:= C_{ijk\ell mp},\end{aligned} \tag{22}$$

while any expectation involving an odd number of spins vanishes. Finally, we note that for any function of the spins $f(\boldsymbol{\sigma})$, we have

$$\frac{\partial}{\partial\beta}\langle f(\boldsymbol{\sigma})\rangle = -\langle f(\boldsymbol{\sigma})E(\boldsymbol{\sigma})\rangle + \langle f(\boldsymbol{\sigma})\rangle\langle E(\boldsymbol{\sigma})\rangle, \tag{23}$$

where $E(\boldsymbol{\sigma}) = -\frac{1}{2}\sum_{ij} J_{ij}\sigma_i\sigma_j - \sum_i h_i\sigma_i$ is the energy of configuration $\boldsymbol{\sigma}$.

We proceed with the the first-order term in Eq. (4). We recall that $\langle\sigma_i\rangle_{\beta=0} = 0$ for all $i \in N$ and we also have

$$\begin{aligned}\langle E(\boldsymbol{\sigma})\rangle_{\beta=0} &= -\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\left(\frac{1}{2}\sum_{ij} J_{ij}\sigma_i\sigma_j + \sum_i h_i\sigma_i\right)\\ &= -\frac{1}{2}\sum_{ij} J_{ij}\delta_{ij} = 0,\end{aligned} \tag{24}$$

since $J_{ii} = 0$ for all $i \in N$ by definition. Thus, the first-order term in Eq. (4) takes the form

$$\begin{aligned}\frac{\partial}{\partial\beta}\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0} &= \Big[-\langle\sigma_i\sigma_j E(\boldsymbol{\sigma})\rangle + \cancel{\langle\sigma_i\sigma_j\rangle\langle E(\boldsymbol{\sigma})\rangle}\\ &\quad \cancel{- \langle\sigma_i\rangle\left(-\langle\sigma_j E(\boldsymbol{\sigma})\rangle + \langle\sigma_j\rangle\langle E(\boldsymbol{\sigma})\rangle\right)}\\ &\quad \cancel{-\langle\sigma_j\rangle\left(-\langle\sigma_i E(\boldsymbol{\sigma})\rangle + \langle\sigma_i\rangle\langle E(\boldsymbol{\sigma})\rangle\right)}\Big]_{\beta=0}\\ &= \frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\sigma_i\sigma_j\left(\frac{1}{2}\sum_{k\ell} J_{k\ell}\sigma_k\sigma_\ell + \sum_k h_k\sigma_i\right)\\ &= \frac{1}{2}\sum_{k\ell} J_{k\ell}\left(\delta_{ij}\delta_{k\ell} + \delta_{ik}\delta_{j\ell} + \delta_{i\ell}\delta_{jk} - 2\delta_{ijk\ell}\right)\\ &= \frac{1}{2}\left(J_{ij} + J_{ji}\right) = J_{ij},\end{aligned} \tag{25}$$

where the penultimate equality follows from J being zero on the diagonal and the final equality follows since J is symmetric.

We now derive the second-order term in Eq. (4). Including only the non-vanishing terms, we have

$$\begin{aligned}\frac{\partial^2}{\partial\beta^2}\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0} &= \left[\left\langle\sigma_i\sigma_j E^2(\boldsymbol{\sigma})\right\rangle - \langle\sigma_i\sigma_j\rangle\left\langle E^2(\boldsymbol{\sigma})\right\rangle\right.\\ &\quad \left.-2\langle\sigma_i E(\boldsymbol{\sigma})\rangle\langle\sigma_j E(\boldsymbol{\sigma})\rangle\right]_{\beta=0}.\end{aligned} \tag{26}$$

Evaluating the second term in Eq. (26), we recall $\langle\sigma_i\sigma_j\rangle_{\beta=0} = \delta_{ij}$ and we have

$$\begin{aligned}\left\langle E^2(\boldsymbol{\sigma})\right\rangle_{\beta=0} &= \frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\left(\frac{1}{2}\sum_{k\ell} J_{k\ell}\sigma_k\sigma_\ell + \sum_k h_k\sigma_k\right)\\ &\quad\cdot\left(\frac{1}{2}\sum_{mp} J_{mp}\sigma_m\sigma_p + \sum_m h_m\sigma_m\right)\\ &= \frac{1}{4}\sum_{k\ell mp} J_{k\ell}J_{mp}C_{k\ell mp} + \sum_{km} h_k h_m\delta_{km}\\ &= \frac{1}{2}\sum_{k\ell} J_{k\ell}^2 + \sum_k h_k^2.\end{aligned} \tag{27}$$

Furthermore, the third term in Eq. (26) is given by

$$\begin{aligned}\langle\sigma_i E(\boldsymbol{\sigma})\rangle_{\beta=0} &= -\frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\sigma_i\left(\frac{1}{2}\sum_{k\ell} J_{k\ell}\sigma_k\sigma_\ell + \sum_k h_k\sigma_k\right)\\ &= -\sum_k h_k\delta_{ik} = -h_i.\end{aligned} \tag{28}$$

Finally, we evaluate the first term in Eq. (26):

$$\begin{aligned}\left\langle \sigma_i\sigma_j E^2(\boldsymbol{\sigma})\right\rangle_{\beta=0} &= \frac{1}{2^n}\sum_{\{\boldsymbol{\sigma}\}}\sigma_i\sigma_j\left(\frac{1}{2}\sum_{k\ell}J_{k\ell}\sigma_k\sigma_\ell+\sum_k h_k\sigma_k\right)\\ &\qquad\cdot\left(\frac{1}{2}\sum_{mp}J_{mp}\sigma_m\sigma_p+\sum_m h_m\sigma_m\right)\\ &= \frac{1}{4}\sum_{k\ell mp}J_{k\ell}J_{mp}C_{ijk\ell mp}+\sum_{km}h_kh_mC_{ijkm}\\ &= \frac{1}{2}\delta_{ij}\sum_{k\ell}J_{k\ell}^2+2\sum_k J_{ik}J_{kj}-2\delta_{ij}\sum_k J_{jk}^2\\ &\quad+\delta_{ij}\sum_k h_k^2+2h_ih_j-2\delta_{ij}h_i^2.\end{aligned} \tag{29}$$

Combining all of the terms in Eq. (26) and canceling appropriately, we are left with

$$\frac{\partial^2}{\partial\beta^2}\left(\frac{1}{\beta}\chi_{ij}\right)_{\beta=0}=2\left(\sum_k J_{ik}J_{kj}-\delta_{ij}\sum_k J_{jk}^2-\delta_{ij}h_i^2\right). \tag{30}$$

All together, Eqs. (5), (25), and (30) represent the zeroth-, first-, and second-order terms in the expansion of $\frac{1}{\beta}\chi_{ij}$ in Eq. (4), respectively. Multiplying by β yields the high-temperature expansion of χ_{ij} in Eq. (6), as desired.

APPENDIX B: LOW-T OPTIMAL EXTERNAL FIELDS FOR GENERAL SYSTEMS

In this appendix, we provide an analytic description of the low-T optimal external fields for a general Ising system, described by J and $\boldsymbol{h}^{(0)}$, and a general external field budget H, generalizing the results of Section IV. For some external field $\boldsymbol{h}\in\mathcal{F}_H$, let $\Omega^0_{\boldsymbol{h}}=\arg\min_{\{\boldsymbol{\sigma}\}}E_{\boldsymbol{h}}(\boldsymbol{\sigma})$ denote the set of ground state configurations under the external field $\boldsymbol{h}^{(0)}+\boldsymbol{h}$, where $E_{\boldsymbol{h}}(\cdot)$ denotes the energy under the external field $\boldsymbol{h}^{(0)}+\boldsymbol{h}$, and let $\Omega^0=\cup_{\boldsymbol{h}\in\mathcal{F}_H}\Omega^0_{\boldsymbol{h}}$ denote the set of all possible ground states that can be induced by some $\boldsymbol{h}\in\mathcal{F}_H$ for an external field budget H.

For sufficiently low temperatures, the ground state of the system dominates any expectation. Thus, any optimal external field will necessarily induce a ground state that has the largest magnetization among the configurations in Ω^0. So, let $\Omega^{0^*}\subset\Omega^0$ denote the subset of possible ground states with the maximum magnetization and let $\mathcal{F}_H^*=\left\{\boldsymbol{h}\in\mathcal{F}_H:\Omega^0_{\boldsymbol{h}}\subset\Omega^{0^*}\right\}$ denote the set of feasible external fields that induce ground states of maximum magnetization.

For sufficiently low temperatures, any optimal external field $\boldsymbol{h}^*$ is located in $\mathcal{F}_H^*$. To differentiate between the external fields in $\mathcal{F}_H^*$, we consider the possible first-excited states, as in Section IV. Let $\Omega^1_{\boldsymbol{h}}$ denote the set of first excited states of the system under the external field $\boldsymbol{h}^{(0)}+\boldsymbol{h}$, and let $\Omega^1=\cup_{\boldsymbol{h}\in\mathcal{F}_H^*}\Omega^1_{\boldsymbol{h}}$ denote the set of possible first-excited states that can be induced by some $\boldsymbol{h}\in\mathcal{F}_H^*$. For $\boldsymbol{h}\in\mathcal{F}_H$, letting $\Delta E(\boldsymbol{h})=\min_{\boldsymbol{\sigma}\in\Omega^1}\max_{\boldsymbol{\sigma}^0\in\Omega^{0^*}}\left(E_{\boldsymbol{h}}(\boldsymbol{\sigma})-E_{\boldsymbol{h}}(\boldsymbol{\sigma}^0)\right)$ denote the energy gap between the ground and first-excited states under the external field $\boldsymbol{h}^{(0)}+\boldsymbol{h}$ and letting $M_0=\sum_i\sigma_i$, for any $\boldsymbol{\sigma}\in\Omega^{0^*}$, denote the magnetization of the optimal ground states, for sufficiently low temperatures, the magnetization takes the form:

$$\begin{aligned}M(\boldsymbol{h}^{(0)}+\boldsymbol{h}) &\approx \frac{\left|\Omega^0_{\boldsymbol{h}}\right|M_0+\left(\sum_{\boldsymbol{\sigma}\in\Omega^1_{\boldsymbol{h}}}\sum_i\sigma_i\right)e^{-\beta\Delta E(\boldsymbol{h})}}{|\Omega^0_{\boldsymbol{h}}|+|\Omega^1_{\boldsymbol{h}}|\,e^{-\beta\Delta E(\boldsymbol{h})}}\\ &\approx M_0+c(\boldsymbol{h})e^{-\beta\Delta E(\boldsymbol{h})},\end{aligned} \tag{31}$$

where $c(\boldsymbol{h})=\frac{1}{\left|\Omega^0_{\boldsymbol{h}}\right|}\left[\left(\sum_{\boldsymbol{\sigma}\in\Omega^1_{\boldsymbol{h}}}\sum_i\sigma_i\right)-\left|\Omega^1_{\boldsymbol{h}}\right|M_0\right]$ is a scalar that depends on $\boldsymbol{h}$. Thus, the low-temperature optimal external fields, i.e., the external fields which maximize the low-temperature magnetization, take the form:

$$\begin{aligned}\boldsymbol{h}^* &= \arg\max_{\boldsymbol{h}\in\mathcal{F}_H^*}c(\boldsymbol{h})e^{-\frac{1}{T}\Delta E(\boldsymbol{h})}\\ &\equiv \arg\max_{\boldsymbol{h}\in\mathcal{F}_H^*}-\operatorname{sign}\left[c(\boldsymbol{h})\right]\Delta E(\boldsymbol{h})\\ &\equiv \arg\max_{\boldsymbol{h}\in\mathcal{F}_H^*}\min_{\boldsymbol{\sigma}\in\Omega^1}\max_{\boldsymbol{\sigma}^0\in\Omega^{0^*}}-\operatorname{sign}\left[c(\boldsymbol{h})\right]\left(E_{\boldsymbol{h}}(\boldsymbol{\sigma})-E_{\boldsymbol{h}}(\boldsymbol{\sigma}^0)\right)\end{aligned} \tag{32}$$

We first note that Eq. (32) is not linear in $\boldsymbol{h}$, and hence the useful algorithmic properties of Eq. (15) have been lost in generalizing to general Ising systems. However, despite being non-linear in $\boldsymbol{h}$, Eq. (32) does reveal a major insight into the structure of low-T optimal external fields in general Ising systems. Depending on the sign of $c(\boldsymbol{h})$, the optimal external field $\boldsymbol{h}^*$ will either maximize or minimize the energy gap. Since the first excited states in $\Omega^1_{\boldsymbol{h}}$ are likely ground states under other external fields (i.e., it is likely that $\Omega^1_{\boldsymbol{h}}\subset\Omega^0$ for $\boldsymbol{h}\in\mathcal{F}_H^*$) and since M_0 is the largest magnetization among the states in Ω^0, we should expect in most cases that $c(\boldsymbol{h})<0$ for all $\boldsymbol{h}\in\mathcal{F}_H^*$. In this case, the low-$T$ optimal external field $\boldsymbol{h}^*$ maximizes the energy gap $\Delta E(\boldsymbol{h})$, focusing H on the loosely-connected nodes, i.e., nodes with opposite parity between the ground and first-excited states. Thus, once we restrict our attention to external fields that induce ground states with the maximum magnetization (i.e., once we restrict to $\boldsymbol{h}\in\mathcal{F}_H^*$), much of the intuition developed in Section IV generalizes naturally to general Ising systems.

[1] R. Albert and A. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.

[2] F. Barahona. On the computational complexity of ising spin glass models. *J. of Physics A: Mathematical and General*, 15(10):3241, 1982.

[3] S. Bharathi, D. Kempe, and M. Salek. Competitive influence maximization in social networks. *International Workshop on Web and Internet Economics*, pages 306–311, 2007.

[4] L. Blume. The statistical mechanics of strategic interaction. *Games and Economic Behavior*, 5:387–424, 1993.

[5] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23 (11):1222–1239, 2001.

[6] C. Castellano, S. Fortunato, and V. Loreto. Statistical physics of social dynamics. *Rev. Mod. Phys.*, 81:591–646, 2009.

[7] E. Demaine, M. Hajiaghayi, H. Mahini, D. Malec, S. Raghavan, A. Sawant, and M. Zadimoghadam. How to influence people with partial incentives. *WWW'14*, pages 937–948, 2014.

[8] P. Domingos and M. Richardson. Mining the network value of customers. *KDD'01. ACM*, pages 57–66, 2001.

[9] N. Du, L. Song, and M. Gomez-Rodriguez. Scalable influence estimation in continuous-time diffusion networks. *NIPS*, 2013.

[10] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. Correlation inequalities on some partially ordered sets. *Comm. in Math. Phys.*, 22(2):89–103, 1971.

[11] S. Galam. Sociophysics: a review of galam models. *Int. J. Mod. Phys. C*, 19(3):409–440, 2008.

[12] S. Galam and S. Moscovici. Towards a theory of collective phenomena: Consensus and attitude changes in groups. *European Journal of Social Psychology*, 21:49–74, 1991.

[13] S. Galam, Y. Gefen, and Y. Shapir. A mean behavior model for the process of strike. *Mathematical Journal of Sociology*, 9:1–13, 1982.

[14] M. Gomez-Rodriguez and S. Bernhard. Influence maximization in continuous time diffusion networks. *ICML*, 2012.

[15] S. Goyal, H. Heidari, and M. Kearns. Competitive contagion in networks. *Games and Economic Behavior*, 2014.

[16] R. Griffiths, C. Hurst, and S. Sherman. Concavity of magnetization of an ising ferromagnet in a positive external field. *J. Math. Phys.*, 11:790, 1970.

[17] D. Isenberg. Group polarization: A critical review and meta-analysis. *J. of Personality and Social Psychology*, 50:1141–1151, 1986.

[18] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

[19] D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. *KDD'03. ACM*, pages 137–146, 2003.

[20] V. Kolmogorov and R. Zabin. What energy functions can be minimized via graph cuts? *IEEE transactions on pattern analysis and machine intelligence*, 26(2):147–159, 2004.

[21] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

[22] D. Levine. Glauber dynamics for ising model. AMS Short Course, 2010. Tutorials.

[23] S. Liu, L. Ying, and S. Shakkottai. Influence maximization in social networks: An ising-model-based approach. *Communication, Control, and Computing (Allerton)*, pages 570–576, 2010.

[24] C. Lynn and D. D. Lee. Maximizing influence in an ising network: A mean-field optimal solution. *NIPS*, 2016 (https://arxiv.org/abs/1608.06850).

[25] M. Mäs, A. Flache, and D. Helbing. Individualization as driving force of clustering phenomena in humans. *PLoS Comput Biol*, 6(10), 2010.

[26] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 7:6–38, 1995.

[27] A. Montanari and A. Saberi. The spread of innovations in social networks. *PNAS*, 107(47).

[28] E. Mossel and S. Roch. On the submodularity of influence in social networks. In *STOC'07*, pages 128–134. ACM, 2007.

[29] M. Moussaïd, J. Kämmer, P. Analytis, and H. Neth. Social influence and the collective dynamics of opinion formation. *PLoS One*, 8(11), 2013.

[30] M. Newman. The structure of scientific collaboration networks. *Proc. Natl. Acad. Sci.*, 98, 2001.

[31] M. Newman and G. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, New York, 1999.

[32] H. Nishimori. *Statistical physics of spin glasses and information processing: an introuction*, volume 111. Clarendon Press, 2001.

[33] H. Nishimori and K. M. Wong. Statistical mechanics of image restoration and error-correcting codes. *Physical Review E*, 60(1), 1999.

[34] A. Nowak, J. Szamrej, and B. Latané. From private attitude to public opinion: A dynamic theory of social impact. *Psychological Review*, 97(3), 1990.

[35] M. Opper and D. Saad. *Advanced mean field methods: Theory and practice*. MIT press, 2001.

[36] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. *KDD'02. ACM*, pages 61–70, 2002.

[37] L. K. Saul, T. Jaakkola, and M. I. Jordan. Mean field theory for sigmoid belief networks. *Journal of artificial intelligence research*, 4(1):61–76, 1996.

[38] T. Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.

[39] K. Sznajd-Weron and J. Sznajd. Opinion evolution in closed community. *Int. Mod. Phys. C*, 11(06).

[40] T. Tanaka. Mean-field theory of boltzmann machine learning. *Phys. Rev. E*, pages 2302–2310, 1998.

[41] M. Teboulle. First order algorithms for convex minimization. IPAM, 2010. Tutorials.

[42] J. Yedidia. *Advanced mean field methods: Theory and practice*. 2001.